Climate Matters, chapter 7
by John Broome (2012)

Chapter 7
Uncertainty

Predictions of the future climate are very uncertain; they range from the relatively benign to the catastrophically harsh. Consequently, uncertainty permeates all the ethics of climate change. We have to take action in response to climate change at a time when we are very unsure what its effects will be. We are equally unsure what will be the effects of anything we do in response to climate change. How should we act in the face of all this uncertainty?

We should not despair at the difficulty. Fortunately, there is a very well-founded theory of how we should act under uncertainty, known as "expected value theory." It applies to uncertainty of any sort, even to the pervasive uncertainty that climate change imposes on us.

It is not the only approach to uncertainty that has been recommended for climate change. Alternatives include the "precautionary principle," which is more a collection of related ideas than a single principle. Some versions of the precautionary principle are demonstrably mistaken. This chapter starts with a review of these and other mistaken approaches to uncertainty, before going on to explain expected value theory.

HOW NOT TO COPE WITH UNCERTAINTY

One wrong way to act in the face of uncertainty is to do nothing. The administration of George W. Bush in the US adopted this policy toward climate change. It put heavy emphasis on the uncertainty of predictions about climate change; indeed, it deliberately exaggerated the uncertainty.[1] Then it presented this uncertainty as a reason for doing nothing to limit emissions of greenhouse gas.

Bush himself explained the policy as early as 2000, before he became president. He said, "There's a lot of differing opinions and before we react I think it's best to have the full accounting, full understanding of what's taking place."[2] In many contexts, this would be a sensible remark. A decision that is based on full information is better than one that is not. If you can costlessly delay a decision till all the information is in, you should delay it. But when delay itself is risky, it is not a sensible remark.

You do not yet know whether your house will catch fire, but that is not a good reason to delay buying a fire extinguisher. By the time you do know, it will be too late. We shall never have full accounting or full understanding of climate change, at least not until it is too late. There is already a strong likelihood that climate change will be damaging, so we

[1] Union of Concerned Scientists, *Scientific Integrity in Policymaking: An Investigation into the Bush Administration's Misuse of Science*, 2004.

[2] Presidential debate, October 11, 2000.

should take precautions against it. Possibly, we may later find the precautions were not needed, but that does not mean we were wrong to take them. If your house never catches fire, that does not mean you were wrong to buy a fire extinguisher.

The Rio Declaration of 1993 condemned the policy of doing nothing long before Bush adopted it. Principle 15 of the declaration states that:

> In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.

This is a version of what is called the precautionary principle. It is a statement of good common sense.

Other versions of the precautionary principle are less sensible. For example, the World Charter for Nature, adopted by the General Assembly of the United Nations in 1982, states that:

> Activities which are likely to pose a significant risk to nature shall be preceded by an exhaustive examination; their proponents shall demonstrate that expected benefits outweigh potential damage to nature, and where potential adverse effects are not fully understood, the activities should not proceed.

The last clause of this statement is the Bush policy in reverse, and just as wrong. It says we should not start an activity until it is known to be beneficial, whereas Bush would not stop an activity until it was known to be harmful. Sometimes we have to take some risks, even when we do not fully understand the potential adverse effects of what we do.

Doing nothing is not the right way to handle uncertainty. An alternative response is to act on the basis of what is likely to happen. More precisely, the idea is that, of the actions available to you, you ought to choose the one that is most likely to produce the best result. This sounds plausible at first, but it too is wrong. Take the fire extinguisher again. It is unlikely that your house will catch fire. So if you do not buy a fire extinguisher, it is most likely that no harm will result, and you will save yourself the cost of a fire extinguisher. On the other hand, if you buy one, it is most likely not to be used, and you will be worse off to the extent of its cost. So if you go by what is most likely to happen, you will not buy a fire extinguisher. But that may well be the wrong choice.

We learn from this example that what is most likely to happen is not necessarily the most important consideration in making a decision. An unlikely possibility may be more important if its results will be extremely bad. It is unlikely but possible that your house will catch fire. The outcome in that case will be extremely bad if you have no fire extinguisher. This may be the most important consideration, so that therefore you ought to buy a fire extinguisher.

MAXIMIZING EXPECTED VALUE

Those are some wrong ways to cope with uncertainty. What is the right way? At heart it accords well with common sense. However, I have to warn you that, when worked out properly, it involves a mathematical concept you may find unfamiliar. The morality of climate change is a quantitative matter—we need to know just what is the size of the response morality calls on us to make. We cannot avoid a little arithmetic in working it out.

Maximizing the expectation of money. A good way to start on the theory of uncertainty is through some examples taken from gambling games. My examples in this chapter are mostly about how to do yourself some good, rather than about morality. That is because the right approach to uncertainty is the same whether you are pursuing your own good or the good of the world. It is easiest to present it using small-scale examples of self-interest. We have some well-tuned intuitions about gambling games, and it is useful to be able to call on them. If you dislike gambling, I ask you to suspend your dislike for the sake of this exercise.

Suppose you are offered a bet on the roll of a single fair dice. You may lay down a stake, and then the dice is rolled. If you take this bet, you will be paid the number that comes up on the dice, in money: $1 if one comes up, $2 if two comes up, and so on. What should you be willing to pay as a stake in order to play this game?

Common sense suggests the answer. We need to work out your average winnings. This is, in dollars, the average of the numbers one to six. This is 3.5. If you were to pay a stake of $3.50, and play this game many times, on average you would break even. Common sense suggests you should be willing to stake anything up to $3.50.

Now suppose you are offered a betting game where the payouts are different. You will get $3 if either one or two comes up on the dice, $4 if three, four, or five comes up, and $12 if six comes up. What stake should you be willing to pay this time? Again, we must calculate the average payout, but this time we need to calculate a weighted average. You have a one-in-three, or 1/3, chance of getting $3, a 1/2 chance of getting $4, and a 1/6 chance of getting $12. To calculate your average payout, each of the different payouts you might receive must be weighted by your chance of receiving it. We multiply $3 by your chance of getting $3, which is 1/3, to get $1. We multiply $4 by your chance of getting it, which is 1/2, to get $2. We multiply $12 by 1/6, to get $2. We add up these three products, to get $5 altogether. This is your weighted average payout. Common sense suggests you should be willing to pay a stake of up to $5 to play the game.

In mathematical terminology, the weighted average I have just described is called an “expectation.” We say that the expectation of your payout, or your “expected payout,” is $5. The terms “expectation” and “expected” are so well-established in the mathematics of uncertainty that I cannot avoid using them. But please remember they are technical terms, and do not mean what those words mean in common English. In the game I described, one payout you definitely do not expect in the common English sense is $5; for sure you will not get $5. Yet this is the expectation of your payout.

When the quantitative outcome of some process is uncertain, the expectation of the outcome is calculated as follows. Take each of the possible values of the outcome and multiply each by the probability of its occurring. Add up all of these products. The sum is the expectation. It is just a weighted average outcome, where the weights are the probabilities.

In the gambling examples, common sense suggests that the right action to take is dictated by expectations. When you have a choice of things to do, you should do the one that gives you the greatest expectation of money. In the games, your choice is whether to play the game or not. You should play if your expected winnings are greater than the stake; otherwise you should not play. To put it briefly, you should maximize your expectation of money. That is the common-sense rule for gambling games.

But is it correct? We should not accept it without question. One argument in support of the common-sense rule is this: Suppose you are offered games like these many times in quick succession. It can be proved mathematically that, if you maximize your expectation of money in each one, you will almost certainly come out ahead in the long run. So you should adopt this policy, and your near certainty of winning in the long run explains why.

This argument is unconvincing. For one thing, you may be offered just one opportunity to play, rather than a series of opportunities. The argument is supposed to show you should maximize your expectation of money in each game, when you have a long sequence of games to play. But it gives no reason to think you should do the same when you have only one game.

For another thing, even when the games are repeated often, the argument begs a question. Although it is almost certain you will win in the long run, it is not completely certain. Even a long sequence of games may go against you. Why should you choose the policy that is almost, but not completely, certain to win? This is another question about how you should act in the face of uncertainty. And we know already from the example of the fire extinguisher that the answer is not necessarily the one the argument assumes. It is not true in all circumstances that you should choose an option that is almost certain to win. You should not choose it if the other possible result, which would happen in the very unlikely case that you did not win, would be utterly disastrous. So the argument as it stands is definitely inadequate.

In any case, we can see that the argument cannot be a good one, because its conclusion is not always true. Maximizing the expectation of money is not always the right way to act in the face of uncertainty.

Maximizing the expectation of value. Insurance provides an example. Suppose you are moving all your possessions from one house to another. Your possessions may be destroyed or stolen on the way. Suppose their value is $50,000, and the chance they will be lost is 1/1000. Then your expectation of loss is $50. However, if you take out full insurance (and if you attach no value to your possessions apart from their money value), you can be sure of losing nothing. Should you take out insurance?

It depends on the premium, of course. Suppose an insurance company offers to insure you for $51. Should you take out insurance for that premium? You will not do so if you follow the policy of maximizing your expectation of money. Your expected loss is $50 if you do not take out insurance, whereas your certain loss from paying the premium is greater, $51. Nevertheless, common sense suggests that in this case you should take out the insurance. Evidently, although common sense supports maximizing the expectation of money for gambling games, it takes a different view for insurance.

Why? Crudely, it is because you should avoid big risks. You should be "risk-averse," as economists say. Why is that? Because $50,000 is big enough to make a noticeable difference to your wealth. That matters because of what economists call the "diminishing marginal benefit of money." This is an important notion in the ethics of risk, and it needs explaining. It will come up later in other contexts too.

The word "marginal" is economist-speak for a small change. In this case it refers to a small change in your wealth. By "benefit" I mean an increase in your well-being, and by "well-being" I mean how well off you are, or how well your life goes. The marginal benefit of money is the amount by which a small increase in your wealth increases your well-being. More precisely, it is the rate at which your well-being increases as you get richer. This marginal benefit normally diminishes as you get richer: the richer you are the less you benefit from further additions to your wealth. If you are poor, a few dollars will buy you some of the necessities of life. If you are rich, you already have all the necessities; all a few dollars will get you is extra trinkets. The first bathroom in your house changes your life; the second not so much. This is the diminishing marginal benefit of money.

In trivial gambling games, common sense tells us you should maximize the expectation of money, but that is because these games make no significant difference to your overall wealth. When you lose a small amount of money, the harm you suffer is proportional to the amount you lose. Losing $2 is twice as bad as losing $1. But when we are dealing with large amounts of money, the diminishing marginal benefit of money kicks in. When the risks are big, benefits and harms are not proportional to gains and losses of money. Because of the diminishing marginal benefit of money, losing a lot of money is proportionately worse for you than losing a little. Losing $50,000 would be more than a thousand times as bad for you as losing $50.

Gains and losses of money matter to you only because of the benefits and harms—gains and losses of well-being—they bring. Money in itself is no good to you; its value to you is the well-being it can buy you. So it seems plausible that you should maximize the expectation of your well-being, rather than the expectation of your money. This was the view taken by Daniel Bernoulli, a mathematician who founded the theory of decision-making under uncertainty in the eighteenth century.[3]

Because losing $50,000 would be more than a thousand times as bad for you as losing $50, the expected harm to your well-being caused by a 1/1000 chance of losing $50,000 is

[3] Daniel Bernoulli, "Specimen theoriae novae de mensura sortis," *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5 (1738), translated by Louise Sommer as "Exposition of a new theory on the measurement of risk," *Econometrica* 22 (1954): 23–36.

greater than the harm caused by losing $50 for sure. We may plausibly assume it is also greater than the harm caused by paying a premium of $51. You should therefore be willing to pay $51 to insure yourself against a 1/1000 chance of losing $50,000. This gives us a plausible explanation of why common sense says you should, in my example, insure your possessions for $51.

We have now arrived at a more general rule for how to act in the face of uncertainty: maximize the expectation of your well-being. The same rule extends to moral contexts. Suppose we have some quantitative notion of how good the world is. When you have to act, and you are uncertain what the results of your act will be, maximize the expectation of the goodness of the world. I call this "expected value theory." I have given no justification for it except that it is supported by common sense. However, the mathematical theory of decision-making under uncertainty does supply a solid justification, which is too complicated to set out in this book.[4] I believe expected value theory to be the correct theory of how we should take uncertainty into account in our moral actions.

Expected value theory tells us how, in principle, we should approach the task of judging the badness of climate change and the goodness of whatever steps we might take to mitigate it. We should value each action we might take—including doing nothing—according to its expected value. That is to say, in principle we should first identify all the different ways the world might go as a result of the action, and then calculate the expected value of all these ways together.

For instance, suppose we hold the concentration of greenhouse gas at 550 ppm. As a result, the world might warm by two degrees eventually and take three hundred years to get there, or by five degrees eventually and take four hundred years to get there, and so on. We judge the probability and the value of each of the possible results. For each, we calculate the arithmetical product of its value and its probability. Then we add up all these products. This gives us the expected value of holding the concentration at 550 ppm. This expected value is what we should care about. In so far as we aim to promote goodness rather than justice, we should aim to maximize expected value.

Probabilities. You will already have spotted a problem with that example. To calculate an expected value, we need to know both the probability and value of each of the possible results, such as a two-degree warming over three hundred years. In practice, we do not have that knowledge. So what should we do? When we try to apply expected value theory to climate change, we have to answer this practical question.

[4] The beginnings of it are in a paper by the philosopher and economist Frank Ramsey, "Truth andprobability," in his *Foundations of Mathematics and Other Logical Essays*, edited by R. B. Braithwaite (London: Routledge and Kegan Paul, 1931). My book *Weighing Goods* reviews the subsequent developments. The mathematics support expected value theory only if we make one contentious assumption, which I call "Bernoulli's hypothesis." *Weighing Goods* examines Bernoulli's hypothesis at length, and declines to support it. However, my later book *Weighing Lives*, ch. 5, offers a revised opinion that supports it. My argument depends on a particular definition of quantities of well-being.

Expected value theory tells us to do our best. We do not know probabilities and values, but we must try to estimate them as well as we can. No one thinks this will be easy.

Values first. To calculate the value of each possible result, as always we have to weigh good features against bad ones. This means applying cost– benefit analysis to each of the possibilities separately. Each possibility will lead to the world's developing in some particular way. There will be a particular growth or shrinkage of population, and people's well-being will improve or diminish in a particular way. We have to set a value on this development. Many questions will arise. What weight should we give to the well-being of future people compared with our own? How should we evaluate the change of population? Later chapters in this book consider questions of valuation like these.

Probabilities are the issue for this chapter. I introduced expected value theory through examples from gambling games because for them probabilities are known. With a gaming device such as a roulette wheel, each possible outcome has an objective probability of occurring, and we can know what that probability is. A roulette wheel can be spun again and again, and its design ensures that each number will show up with a regular frequency; with a European wheel, in the long run a particular number—say, 15—will show up one time in thirty-seven. This frequency of 1/37 constitutes an objective probability of the number 15. We know what this probability is. When an event can be repeated again and again, the probability of any particular outcome is its frequency, which can be known.

But most of the uncertainty in our lives is not like that. Will you be able to get away from the meeting in time to catch the train at 2:30? Will the car start in the snow tomorrow? Questions like these can rarely be reduced to frequencies that constitute objective probabilities. These are one-off events that are not regularly repeated. For one-off events, probability is different. It is a matter of rationality rather than frequency. The question is: How much credence should you rationally give to the possibility that the car will not start, or to the possibility that you will miss the train? The answer will depend on the evidence you have. You should take into account who is chairing the meeting, or what temperature is forecast tomorrow and how old the car's battery is, and anything else that has a bearing. The probability you assign to your catching the train or to your car's starting should be based on all the evidence you can muster.

The more you can muster, the more tightly the evidence will determine the probability. In the extreme, the evidence will amount to a frequency. If you run each week from the meeting to the train, and you find you catch the train two times out of three, you should attribute a two-thirds probability to catching the train this time. At the other extreme, the evidence may be weak and indefinite, and then it does not fully determine what probability you should assign to the event. In cases of that sort, equally reasonable people may assign different probabilities from each other.

Climate change is a case in which the evidence does not tightly determine what probabilities should be assigned to the various possibilities. However, the situation improves as time passes. Scientists progressively acquire more and more evidence that helps to narrow probabilities, and they pass it on to the rest of us. They are already willing to supply us with a few probabilities. Here is just one small example from the most recent

IPCC report. For each of a number of “scenarios,” which specify how population and the economy develops and how well technology responds to climate change, the report describes a “likely range” of temperature increases over the next century. It says there is a two-thirds probability that the actual increase will be within this range. For instance, in the B2 scenario (whose details we can ignore here), there is a two-thirds probability that the temperature increase will be between 1.4 and 3.8 degrees.

To work out expected values, we need much more detailed information about probabilities than that. While we are waiting for it to be supplied by scientists, we have to do the best we can. On the basis of what evidence we can muster, we have to assign each event a probability as well as we can.

In any case, the lack of firm probabilities is not a reason to give up expected value theory. You might despair and adopt some other way of coping with the uncertainty; you might adopt some version of the precautionary principle, say. That would be a mistake. Stick with expected value theory, since it is very well founded, and do your best with probabilities and values.

HOW SHOULD WE RESPOND TO THE SMALL CHANCE OF CATASTROPHE?

Expected value theory explains why the most likely result of what we do is not necessarily the most important result. An unlikely possibility may be so bad that, even multiplied by the small chance that it will happen, its badness is much more important than what is likely to happen. I gave an example earlier. Your house is not likely to catch fire. Nevertheless, it will be a disaster if it does, and for this reason it may well be that you ought to buy a fire extinguisher.

Climate change may be another example of the same type. There is a real risk it will lead to a terrible disaster. In chapter 2 I defined climate sensitivity as the basic measure of how severe the greenhouse effect is. The latest report of the IPCC assigns a small probability—more than 5 percent—to its being as high as six degrees, and even some probability to its being as high as ten degrees or more.[5] Climate sensitivity is the increase in temperature that would occur in the long run if the concentration of carbon dioxide was effectively doubled from pre-industrial levels, and then held at that doubled level (which is about 550 ppm). It is not directly a prediction of future temperature. However, unless the world takes drastic action soon, concentrations of carbon dioxide will be effectively doubled within a few decades, and will continue to rise after that. We therefore have to reckon on a small probability of six-degree and even ten-degree warming in the long run, unless the world takes drastic action.

These are extreme temperatures. During the last ice age about 20,000 years ago, global temperatures were about five degrees colder than they are now. Five degrees of cooling gave us an ice-age; six degrees of warming must be expected to give us dramatic effects in the other direction. It is hard to form any idea of what ten degrees would bring. The Earth has not been that warm for tens of millions of years. Most of Antarctica would melt, causing

[5] IPCC, *Fourth Assessment Report*, ch. 9.

the sea to rise by anything up to 70 meters. The margins of the continents would drown, and with them most of the world's great cities and vast tracts of farmland.

Ten degrees of warming would be a great catastrophe. It would cause dreadful destruction and suffering. It would also entail a collapse of the planet's human population because, after the loss of so much of the best land and such a radical change in its climate, the Earth could not sustain a population as large as it has now. Given the economic and political disruption that would follow such an extraordinary revision of the Earth's geography, we should not be confident humanity would survive at all. True, it would take many centuries for temperatures to rise to such high levels, and longer for Antarctica to melt. But compared with the history of human beings on Earth, centuries do not constitute a long delay before extinction.

It is a real possibility that climate change could lead to a catastrophe. Expected value theory tells us that, in assessing the badness of climate change, we have to think in terms of expectations. The expectation of harm caused by a catastrophe is the badness of the catastrophe multiplied by the very small probability that it will happen. The view is gaining ground among economists that this possibility is what we should worry about more than anything else.[6] The most likely result of climate change is warming of a few degrees. But the view is that the possible catastrophe of a greater increase would be so bad that, even multiplied by its very small probability, its expected badness is more important than the harm that would be caused by this most likely result.

If this is right, it implies that we ought to do the equivalent of buying a fire extinguisher. We ought to take drastic action to reduce emissions of greenhouse gas. But is it right? We cannot give a responsible answer to that question without taking on two tasks. One is to improve our estimates of what is the true probability of catastrophe. This is a task for science. The other is to judge just how bad the catastrophe would be. How bad would it be for the world's human population to crash? How bad would the extinction of humanity be? This is a task for moral philosophy.

And so we must continue with developing the theory of value within moral philosophy. We can now see that this theory will have to be pushed a long way. The possibility of catastrophe raises extreme questions for the theory of value. They will be difficult to answer, but philosophy does not run away from difficult questions. In chapter 10 I hope to show how we can make a start on answering them; in this book I can go no further than that.

In the meantime, there are some less extreme questions of value still to be tackled. We need to consider how to value future events compared with present ones, and we need to think how to incorporate the value of people's lives into our valuations. Those are the topics of the next two chapters.

---

[6] It is particularly associated with the economist Martin Weitzman. See his paper "On modeling and interpreting the economics of catastrophic climate change," *Review of Economics and Statistics* 91 (2009): 1–19.